arXiv:1606.08395v1 [math.NA] 23 Jun 2016

Polynomial Form of the Matrix Exponential

Daniel Gebremedhin, Charles Weatherford

Physics Department, Florida A&M University, Tallahassee, FL, USA.

Abstract

An algorithm for numerically computing the exponential of a matrix is presented. We have derived a polynomial expansion of e^x by computing it as an initial value problem using a symbolic programming language. This algorithm is shown to be comparable in operation count and convergence with the state–of–the–art method which is based on a Pade approximation of the exponential matrix function. The present polynomial form, however, is more reliable because the evaluation requires only linear combinations of the input matrix. We also show that the technique used to solve the differential equation, when implemented symbolically, leads to a rational as well as a polynomial form of the solution function. The rational form is the well-known diagonal Pade approximation of e^x. The polynomial form, after some rearranging to minimize operation count, will be used to evaluate the exponential of a matrix so as to illustrate its advantages as compared with the Pade form.

Keywords: matrix exponential, polynomial form, sparse matrix

1. Introduction

The exponential function (EF) of a square matrix (matrix exponential function–MEF) is one of the most important functions in a computational linear algebra. It is a fundamental topic of research pertaining to functions of matrices. Since it can be modeled as a solution to an initial value problem, such a technique developed for its successful computation can likely be adopted to solve other problems of physical interest. The solution of the time–dependent Schrödinger equation (TDSE) is such a problem. [1].

Many algorithms have been developed by different authors to numerically compute the MEF [2]. Perhaps the most successful of them is based on the

Email addresses: daniel1.gebremedhin@famu.edu (Daniel Gebremedhin), charles.weatherford@famu.edu (Charles Weatherford)

Pade approximation (PA) of the MEF. A very efficient implementation of the algorithm is published in [3] and is adopted in programming languages such as *MATLAB* [4] and other numerical scientific libraries.

The PA of the EF is a very compact rational expansion of e^x about $x = 0$. When applied to a matrix, its evaluation must eventually involve a matrix inversion, which can sometimes lead to a (nearly) singular system of equations. This issue has been pointed out and discussed in [2] in some detail. Hence, a polynomial representation of the EF, which is as convergent and compact as the PA can be very favorable for implementation on the MEF. In the present work, such a polynomial form has been derived. We will demonstrate that it approximates the MEF with similar convergence as its PA and that it can be computed with comparable efficiency.

In this paper, we will use variables x, τ etc to represent general square matrices. References to scalar quantities should be clear from context. We will use the term norm to mean 1–norm ($\|x\|_1$) of matrix [5]. The relative error of matrices are also calculated based on the 1–norm.

2. Preliminary

The PA can be thought of as an order $2m \to m$ numerical economization of a polynomial representation of a function. Given the coefficients of a general polynomial of order $2m$, it is possible to calculate the corresponding coefficients of two other polynomials of order m so that their quotient results in a rational approximation of the same function. This simple rearrangement usually leads to better convergence, which makes it an interesting topic of numerical studies. Efficient numerical algorithms are available that can calculate these PA coefficients for an arbitrary polynomial [6]. When the same algorithm is applied to a Taylor series expansion of the EF, one can arrive at its diagonal PA. In fact, these two representations of the exponential function have known forms denoted here by T and Q respectively as:

$$e^x \simeq T_{2m}(x) \quad \to \quad e^x \simeq Q_m(x) = \frac{R_m(x)}{R_m(-x)} \tag{1}$$

where,

$$T_m(x) = \sum_{\mu=0}^{m} \frac{x^\mu}{\mu!}, \quad R_m(x) = \sum_{\mu=0}^{m} \binom{2m-\mu}{m} \frac{x^\mu}{\mu!}. \tag{2}$$

While the advantage in reduction of operation count that Q has over T in approximating the EF is immediately apparent from Eq. (1), simple numerical tests reveal that the resulting increase in convergence is also quite significant. This improvement in accuracy exhibited by the PA is generally

function dependent [6]. In the following section, we will describe an algorithm that can be used to generate an approximate representation to the EF as a solution of an initial value problem. By running the algorithm in symbolic programming language we were able to observe that the method can lead to both rational and polynomial expressions of the EF, similar to the forms of Q and T above, if the appropriate set up of input parameters are used. The resulting rational expression happens to be identical to the diagonal PA (Q) given in Eq. (1). Under the appropriate circumstances, a similar algorithm can be employed to derive rational approximations of other elementary functions also. The derived polynomial expression however, will be the main topic of this paper, and will be used to evaluate the MEF. In the process we will reveal some interesting contrasts to the PA, which is the current state–of–the–art method for evaluating the MEF. Finally numerical tests and conclusions will be presented.

3. Description of Algorithm

Let $S \equiv d/dt + p(t)$ be a first order differential operator and $F(t)$ its solution function satisfying the following ordinary differential equation (DE).

$$SF(t) = q(t), \qquad t_1 \leq t \leq t_2 \tag{3}$$

where, t_1 and t_2 represent end points of a particular finite element in t. Let τ be a local variable with domain $-1 \leq \tau \leq 1$ defined by the linear transformation

$$t = \frac{1}{2}\left[(t_2 - t_1)\tau + (t_2 + t_1)\right]. \tag{4}$$

In terms of the local variable τ, Eq. (3) will be re–written as

$$\bar{S}\bar{F}(\tau) = \bar{q}(\tau) \tag{5}$$

where the meaning of the over-bar is clear. At this point we will expand $\bar{F}(\tau)$ in a basis set that let us explicitly fix the initial value of the function at t_1 which corresponds to $\tau = -1$ in terms of the local time.

$$\bar{F}(\tau) = \sum_{\mu=0}^{M-1} s_\mu(\tau) B_\mu + \bar{F}(-1) \tag{6}$$

where s is defined in terms of Legendre polynomials of the first kind (P) as in [7].

$$\begin{aligned} s_\mu(\tau) &= \int_{-1}^{\tau} P_\mu(t)\mathrm{dt} \\ &= \frac{1}{2\mu+1}\left[P_{\mu+1}(\tau) - P_{\mu-1}(\tau)\right] \end{aligned} \tag{7}$$

with $s_0(\tau) = 1+\tau$. The s-functions satisfy the following recurrence relation.[8]

$$s_\mu(\tau) = \frac{1}{\mu+1}\left[(2\mu-1)\tau s_{\mu-1}(\tau) - (\mu-2)s_{\mu-2}(\tau)\right] \tag{8}$$

Note that $s_\mu(-1) = 0$ and derivative of $s_\mu(\tau)$ is $P_\mu(\tau)$.

Substituting the expansion given in Eq. (6) into Eq. (5), projecting from the left by $P_\nu(\tau)$ and integrating over τ results in the following set of simultaneous equations of size M:

$$\sum_{\mu=0}^{M-1} \bar{\Omega}_{\nu\mu} B_\mu = \bar{\Gamma}_\nu \tag{9}$$

where,

$$\begin{aligned} \bar{\Omega}_{\nu\mu} &= \int_{-1}^{1} P_\nu(\tau)\bar{S}s_\mu(\tau)\mathrm{d}\tau \\ \bar{\Gamma}_\nu &= \int_{-1}^{1} P_\nu(\tau)[\bar{q}(\tau) - \bar{p}(\tau)\bar{F}(-1)]\mathrm{d}\tau. \end{aligned} \tag{10}$$

From Eq. (7) it is clear that when $p(t)$ is a constant, $\bar{\Omega}$ is a tridiagonal matrix. After solving Eq. (9) for B, we can evaluate the solution function from Eq. (6). The value at the end point t_2 is particularly important for propagation of the solution and has the following simple form [8]

$$\bar{F}(+1) = 2B_0 + \bar{F}(-1). \tag{11}$$

Note that Eq. (6) is a polynomial of order M in τ as can be seen from the form of s in Eq. (7).

Now, we will set up the parameters in the algorithm such that the resulting solution function is e^x.

3.1. Rational Approximation to the EF

Let $F(t) = e^{tx}$, where $0 \le t \le 1$. The EF can be calculated as a solution to an initial value problem by employing the above algorithm with $p(t) = -x$, $q(t) = 0$, $t_1 = 0$, $t_2 = 1$ and $F(0) = 1$. The desired expression for the EF is obtained at $F(1) = e^x$ from Eq. (11). By letting x to be an undetermined variable, the calculation was done in a symbolic programming language *Mathematica* [9]. For m (number of basis functions), the resulting expression for the EF is found to be identical to $Q_m(x)$ in Eq. (1). It is interesting that one can calculate the coefficients for the PA of the EF directly from the DE in this way. For the sake of completeness, if we break the t axis into, say k, uniform size finite elements, the end result was found to be $F(1) = (Q_m(x/k))^k$. This is consistent with the identity of the EF given by $e^{kx} = (e^x)^k$. This property is more efficiently exploited by the method of scaling and squaring, which chooses k to be a power of 2.

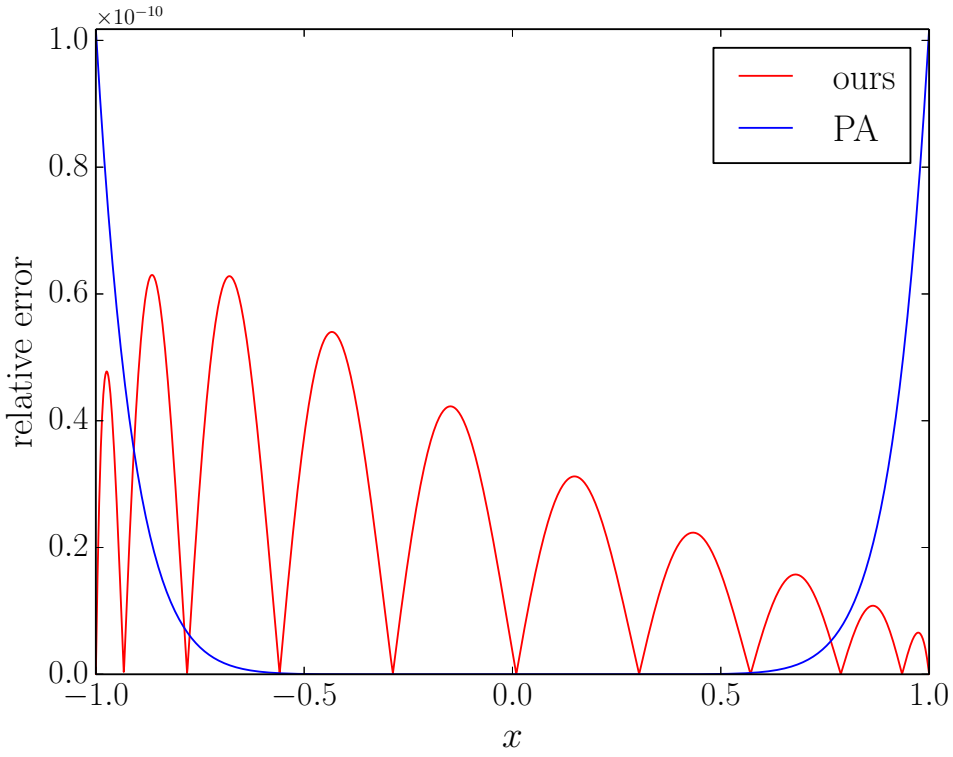

Figure 1: Relative error plots of $E_{10}(x)$ with $\theta = 1$, and $Q_5(x)$ in approximating e^x.

3.2. Polynomial Approximation to EF

Lets now drop x and define $F(t) = e^t$, where $-\theta \leq t \leq \theta$, so that we can use the above algorithm to have a polynomial expression of the EF. θ is a positive number which determines the domain of the result, to be chosen later based on number of basis functions and the working machine precision. The parameters for the algorithm will be set as: $p(t) = -1$, $q(t) = 0$, $t_1 = -\theta$, $t_2 = \theta$ and $F(t_1) = e^{-\theta}$. The output of the algorithm will be in a polynomial form given in Eq. (6) which will be valid for any τ. Specifically, we will denote it by E as shown below.

$$e^x \simeq \bar{F}(\tau = x/\theta) = E_M(x) \tag{12}$$

We anticipate the accuracy of $E_{2m}(x)$ to be comparable with $Q_m(x)$ of PA. Fig. 1 shows a plot of relative errors, in approximating the scalar e^x, of the two methods for $m = 5$. i.e. $E_{10}(x)$ and $Q_5(x)$, where $\theta = 1$. The span θ has been deliberately made too wide because lowering the accuracy close to working precision would have altered the plots to have the familiar random shape. Note the scale of the vertical axis. The expansion for the

PA is centered about the origin, and hence, the U–shape; while the new polynomial is a result of a spectral method which produces the kind of error distribution shown. Within the domains of θs considered in this paper, the upper bound of the relative error for the PA is generally slightly higher than the way it is portrayed in the typical plot. This establishes the fact that $E_M(x)$, although it is a polynomial representation, is much more convergent than the Taylor approximation $T_M(x)$ given in Eq. (1).

In passing, we notice that the choice of $p(t) = i\hat{H}$, $q(t) = 0$ where, $\hat{H}$ is the Hamiltonian matrix and i imaginary number, casts the problem into a form of Schrödinger equation [10, 11, 12]. As will be exhibited soon, its performance in a purely mathematical setting promises a favorable prospect for the algorithm to be applied in physical systems. The techniques discussed in this article can be adopted to effectively propagate solutions to TDSE and will be reported as soon as the details are worked out.

4. Evaluating the Polynomial Function

In this section we will consider how to efficiently evaluate the EF from its polynomial form discussed above. This can be accomplished by minimizing the number of matrix multiplications (MMs) needed for its evaluation. We will also choose the parameters such as the order of expansion M and the corresponding span θ so that, within the given domain, the EF can be calculated close to the working machine precision.

4.1. Product Form

Constructing $E_M(x)$ using the recursion relation shown in Eq. (8) can only be done by $(M-1)$ matrix multiplications. We need to lower the number of MMs by at least a factor of 4 in order to be competitive with present methods for similar accuracy. In [3], the MMs were lowered primarily because both of the polynomials in the quotient of the PA are already of order $M/2$. In this paper, we will rearrange our polynomial, not as quotient, but as a product of other polynomials of smaller order. This can be efficiently done by calculating all the roots $x_1, x_2, \ldots x_M$ of the polynomial given in Eq. (6), after substituting for $\tau = x/\theta$. This allows us to rewrite it as a product of $(x-x_1)(x-x_2)\ldots(x-x_M)$. Generally, roots of a polynomial can be complex numbers, in which case, they always are complex conjugate pairs. So we need to inflate those terms that belong to conjugate pairs, say, x_1^* and x_1, into real quadratic expressions as $(x-x_1)(x-x_1^*) = [x^2 - 2(x_1 + x_1^*)x + x_1^* x_1]$. Hence, using this procedure, we can generally write the EF in the form

$$E_M(x) = \alpha \prod_{i=1}^{M/2} \left(\sum_{j=0}^{2} c_{ij} x^j \right) \tag{13}$$

where all the coefficients c_{ij} are now real numbers. M is assumed to be an even integer. α takes the value of the leading coefficient in Eq. (6), i.e., the coefficient of x^M. Now, by storing x^2 we can construct E by a total of $M/2$ multiplications which is clearly an improvement over $M-1$. We can further consider polynomials of order 4, 6 etc., and seek the most economical arrangement. Finding the order of the polynomial in the above product that requires the least number of MMs is essentially an optimization problem of a general case given by,

$$E_M(x) = \alpha \prod_{i=1}^{m'} \left(\sum_{j=0}^{m} c_{ij} x^j \right) \tag{14}$$

where, $M = m'm$. Storing $x^2, \ldots, x^m$, which takes $(m-1)$ multiplies, allows us to construct any of the polynomials in the above product. Then another $(m'-1)$ multiplications are needed in order to complete the evaluation of E. Hence, the total number of MMs required is given by the function $(M/m + m - 2)$. The minimum of this function occurs at $m = m' = \sqrt{M}$, giving a total of $2(\sqrt{M}-1)$ multiplications. Apparently, convenient values for the order of the polynomial are squares of even integers $M = 2^2, 4^2, 6^2, \ldots, (2m)^2$, respectively requiring $2, 6, 10, \ldots, 2(2m-1)$ multiplications. In this particular choice of M, a unit increase in m always raises the number of MMs by 4.

Note that rearranging a polynomial into a product of other polynomials of lower order, as discussed above, leads to an exactly equivalent expression unlike the rational form of PA which can generally alter (usually for the better) the convergence of the corresponding polynomial form.

Once the orders of the polynomial M to be used have been selected, we need to fix the corresponding θ so that we can calculate the required coefficients c.

4.2. Choice of Span θ

In a recent paper [13], we have defined an adaptive finite element step size choice, which is based on Taylor series expansion, that is effective for solving differential equations. When the method is applied to the EF the result is a step size of 1.5. We will use this quantity as a unit of measure of span and choose from values of $\theta = 0.75m$, $m = 1, 2, \ldots$ By making relative error plots of the scalar function similar to Fig. 1, the largest θ value with error bounds reasonably within the required machine precision has been selected.

Table 1 shows a summary of M and θ values considered in the present work.

Table 1: Values of parameters for different orders M are shown. The 2^{nd} column is the number of multiplications required with comparison from $\pi_{M/2}$ of [3]. θ is the value of span for the corresponding machine precision ϵ. Equivalent values in [3] are displayed as $\theta_{M/2}$

M	$2(\sqrt{M}-1)$	θ	ϵ	$\pi_{M/2}$	$\theta_{M/2}$
16	6	1.5	2^{-52}	5	1.5
36	10	9.75	2^{-52}	8	11.0
64	14	20.25	2^{-52}	–	–
64	14	12	2^{-112}	–	–

We have also included a span for $M = 64$ suitable for calculations in quadruple precision. Its output will be used as an exact value of the EF for comparison purposes. Note that the method of [3] still has to solve the resulting matrix equation after constructing the matrices in the quotient of PA by performing the indicated $\pi_{M/2}$ MMs.

For a given order M, all the parameters $[\alpha, c_{ij}, \theta]$ need to be calculated only once and stored. All of these calculations that are necessary to determine the final set of parameters have been done using exact symbolic calculations in *Mathematica*. In Table 2 the first 19 figures of such a result for $M = 16$ are displayed for demonstration purposes.

5. Numerical Tests

We have extensively tested the results of $E_M(x)$ on different kinds of matrices. We will present the results of three sets of examples herein. In all test cases we will show a comparison with the *expm* function of *MATLAB* which clearly asserts that it implements the PA according to [3]. Following *expm*, no preconditioning of the input matrices such as trace reduction or matrix balancing suggested in [3] has been done for comparison purpose.

Matrices with norms less than 1.5 and 9.75 will be handled by $M = 16$ and 36 respectively, while the ones with higher norms will be calculated by $M = 36$ with proper use of scaling and squaring to lower the norm of the matrices to below $\theta = 9.75$. As mentioned earlier, relative errors will be calculated with reference to the output of $E_M(x)$ in quadruple precision using parameters shown in the last row of Table 1.

When elements of the resulting EF of the matrix overflow beyond 2^{1024} in absolute value, the corresponding relative errors are shown as 1 in the plots. Similarly, relative errors less than 10^{-17} are overwritten to 10^{-17}. The

Table 2: Numerical values of the indicated parameters for $M = 16$. The first 19 figures are shown. All the leading coefficients are unity ($c_{i4} = 1$). The number in square bracket signifies power of 10.

parameter	value
α	4.955887515892002289[-14]
c_{13}	-4.881331340410683266
c_{12}	-14.86233950714664427
c_{11}	862.0738730089864644
c_{10}	3599.994262347704951
c_{23}	7.763092503482958289
c_{22}	77.58934041908401266
c_{21}	430.8068649851425321
c_{20}	1693.461215815646064
c_{33}	9.794888991082968084
c_{32}	98.78409444643527097
c_{31}	387.7896702475912482
c_{30}	1478.920917621023984
c_{43}	3.323349845844756893
c_{42}	37.31797993128430013
c_{41}	545.9089563171489062
c_{40}	2237.981769593417334
θ	1.5

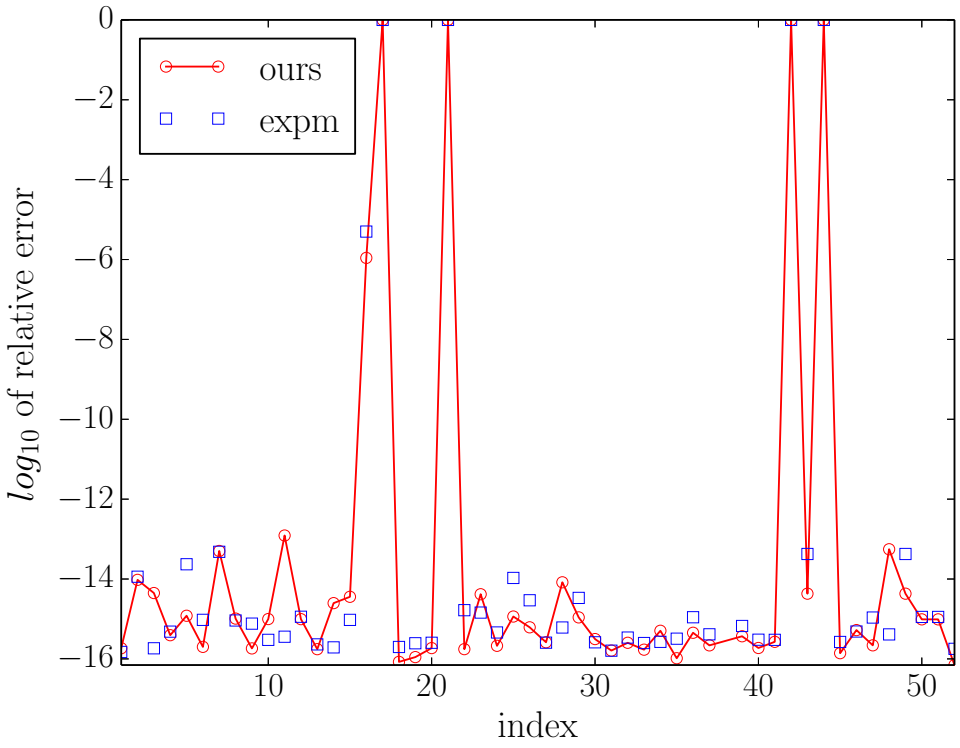

Figure 2: Test results for 51 real matrices of size 8 constructed using the matrix computation toolbox. Log_{10} of the relative errors are shown.

programming has been done in Fortran 95, using GNU gcc version 4.9.1 compiler, in a 2.5 GHz Intel Core i7 MacBook Pro laptop computer.

5.0.1. The Matrix Computation Toolbox

The first test matrices are constructed using the subroutines given in the matrix computation toolbox [1] and [14]. The size of all matrices has been set to be 8×8. Fig. 2 shows the result of the numerical test with indexes adopted from their catalog. Matrices $17, 21, 42, 44$ are beyond the scope of the test due to overflow. On the rest of the matrices, there is a striking qualitative similarity in accuracy between the two functions.

5.0.2. Matrix Market

The second set of test matrices have been taken out of the matrix market in [2] and [15]. A query has been submitted for real and square matrices of size

[1] http://www.maths.manchester.ac.uk/ higham/mctoolbox/
[2] http://math.nist.gov/MatrixMarket/index.html

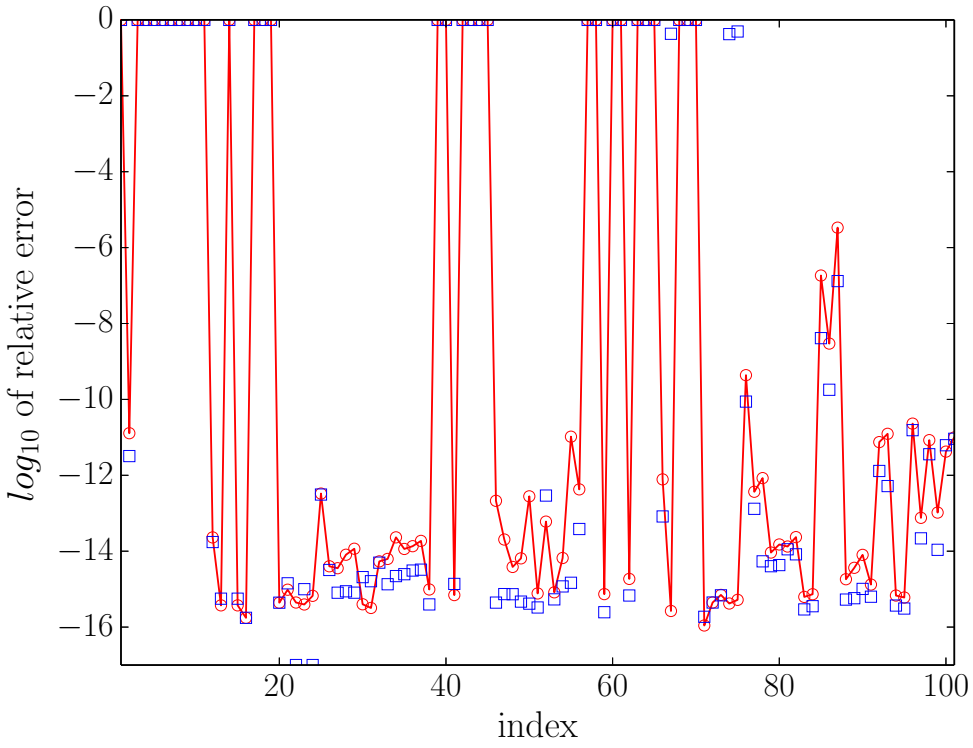

Figure 3: Test results for 101 real matrices of sizes up to 500, downloaded from matrix market. Log_{10} of the relative errors are shown.

up to 500. This returned a set of 101 different kinds of matrices all of which has been tested here. Relative errors of *expm* and ours is shown in Fig. 3. In both functions, 30 matrices overflow upon evaluation. The qualitative similarity in accuracy between the two methods is consistent here as well, except for three matrices labeled with indexes $67, 74, 75$, in which there is a clear difference in accuracy in favor of our method. Those matrices are respectively named 'mhd416b', 'plat362' & 'plskz362' in the matrix market and have sizes $416, 362$ & 362. These three matrices commonly have low norms and high sparsity, which among other reasons, leads us to suspect that they might have posed an ill–conditioned matrix during the matrix inversion step of the PA. To see if sparsity is an issue, we have made a test on collection of matrices which are mainly sparse.

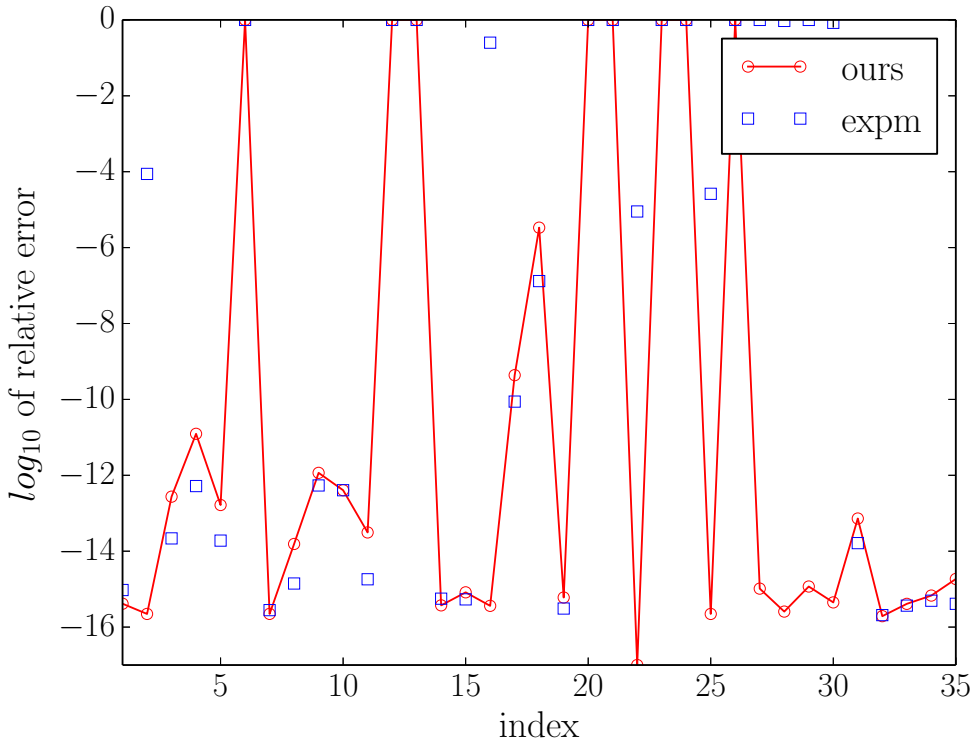

Figure 4: Test results for 35 real matrices of sizes up to 100, downloaded from UF sparse matrix collection. Log_{10} of the relative errors are shown.

5.0.3. UF Sparse Matrix Collection

The last test is downloaded from University of Florida sparse matrix collection [3] [16]. We downloaded and tested all real square matrices with sizes up to 100. The number of such matrices was 35. The result of the test is plotted in Fig. 4. The plot clearly exposes the weakness of the *expm* function in addressing sparse matrices. The matrix labeled by index 27, named 'dbGD97_b' in the collection, for example, completely blows up to *Inf* and/or *NaN*, when evaluated by *expm* which can only be explained by a singular matrix during an LU decomposition process.

It is not clear how to *a priori* identify what kind of input matrix will eventually lead to a poorly conditioned matrix which will compromise the inversion step in the PA. This makes our method more reliable because it merely involves linear combinations of the input matrix. Note that the above tests are exhaustive in the sense that all the resulting matrices that fulfill the

[3] http://www.cise.ufl.edu/research/sparse/matrices/

mentioned search criteria are considered, and there are no cases where *expm* outperforms ours other than what is shown in the plots.

Finally, although not shown here, by changing the horizontal axis in the relative error plots to be the norm of the input matrix (instead of index), we were able to see that there is a compelling correlation between the two. Even with the powerful method of scaling and squaring in place, very high norm of the input matrix is known to be a challenge inherent to computing the EF, which seems to be the case with our algorithm as well.

6. Conclusion

We have derived a numerical algorithm that calculates the EF of matrix. The EF is given in a polynomial form, which we have shown how it can be evaluated by a minimal number of MMs. Sparsity of matrices can be exploited element–wise during evaluation of MMs, which makes our method more so efficient. Matrices with very high norms correspond to longer propagation of solution, which naturally compounds error growth as long as we are working with finite precision. It is those matrices with very high norms that were challenging to our algorithm.

Generally, in this work we have implemented a numerical method that enabled the derivation of both rational and polynomial expansions of an important mathematical function from its DE. This algorithm can readily be adopted to a more realistic dynamical systems that can be modeled as an evolution of initial value problem [13].

The similarity of results shown on the test plots and other aspect of the calculations indicate that our polynomial form is indeed complementary to PA. It is informative to see that they are both solutions to the same DE attainable via a simple technique. But the polynomial form is based solely on MM and avoids the matrix inversion process altogether, and hence, as the plotted test results indicate, ours is the less dubious one.

7. Acknowledgement

DHG and CAW were partially supported by the Department of Energy, National Nuclear Security Administration, under Award Number(s) DE-NA0002630. CAW was also supported in part by the Defense Threat Reduction Agency.

8. REFERENCES

[1] D. Tannor, Introduction to Quantum Mechanics: A Time-dependent Perspective, University Science Books,

Sausalito, California, USA, ISBN 9781891389238, URL http://books.google.com/books?id=t7m08j3Wi9YC, 2007.

[2] C. Moler, C. Van Loan, Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later, SIAM Review 45 (1) (2003) 3–49, doi:10.1137/S00361445024180, URL http://dx.doi.org/10.1137/S00361445024180.

[3] N. Higham, The Scaling and Squaring Method for the Matrix Exponential Revisited, SIAM Journal on Matrix Analysis and Applications 26 (4) (2005) 1179–1193, doi:10.1137/04061101X, URL http://dx.doi.org/10.1137/04061101X.

[4] MATLAB, Version 8.4 (R2014b), The MathWorks Inc., Natick, Massachusetts, 2014.

[5] G. Golub, C. Van Loan, Matrix Computations, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, Maryland, USA, ISBN 9781421407944, URL http://books.google.com/books?id=X5YfsuCWpxMC, 2013.

[6] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, Numerical Recipes 3rd Edition: The Art of Scientific Computing, Cambridge University Press, New York, NY, USA, 3 edn., ISBN 0521880688, 9780521880688, 2007.

[7] F. W. Olver, D. W. Lozier, R. F. Boisvert, C. W. Clark, NIST Handbook of Mathematical Functions, Cambridge University Press, New York, NY, USA, 1st edn., ISBN 0521140633, 9780521140638, 2010.

[8] D. Gebremedhin, Computational Methods in Few–Body Physics, PhD dissertation, Florida A&M University, Department of Physics, a full PHDTHESIS entry, 2013.

[9] Mathematica, Version 7.0 1.0, Wolfram Research, Champaign, Illinois, 2009.

[10] K. Kormann, S. Holmgren, H. O. Karlsson, Accurate time propagation for the Schrdinger equation with an explicitly time-dependent Hamiltonian, The Journal of Chemical Physics 128 (18) 184101, doi:http://dx.doi.org/10.1063/1.2916581, URL http://scitation.aip.org/content/aip/journal/jcp/128/18/10.1063/1.2916581.

[11] H. TalEzer, R. Kosloff, An accurate and efficient scheme for propagating the time dependent Schrdinger equation, The Journal of Chemical Physics 81 (9) (1984) 3967–3971, doi:http://dx.doi.org/10.1063/1.448136, URL http://scitation.aip.org/content/aip/journal/jcp/81/9/10.1063/1.448136.

[12] A. Leclerc, G. Jolicard, D. Viennot, J. P. Killingbeck, Constrained adiabatic trajectory method: A global integrator for explicitly time-dependent Hamiltonians, The Journal of Chemical Physics 136 (1) 014106, doi:http://dx.doi.org/10.1063/1.3673320, URL http://scitation.aip.org/content/aip/journal/jcp/136/1/10.1063/1.3673320.

[13] D. H. Gebremedhin, C. A. Weatherford, Calculations for the one-dimensional soft Coulomb problem and the hard Coulomb limit, Phys. Rev. E 89 (2014) 053319, doi:10.1103/PhysRevE.89.053319, URL http://link.aps.org/doi/10.1103/PhysRevE.89.053319.

[14] N. J. Higham, Accuracy and Stability of Numerical Algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edn., ISBN 0-89871-521-0, 2002.

[15] R. F. Boisvert, R. Pozo, K. Remington, R. F. Barrett, J. J. Dongarra, Matrix Market: A Web Resource for Test Matrix Collections, in: Proceedings of the IFIP TC2/WG2.5 Working Conference on Quality of Numerical Software: Assessment and Enhancement, Chapman & Hall, Ltd., London, UK, UK, ISBN 0-412-80530-8, 125–137, URL http://dl.acm.org/citation.cfm?id=265834.265854, 1997.

[16] T. A. Davis, Y. Hu, The University of Florida Sparse Matrix Collection, ACM Transactions on Mathematical Software 38 (1) (2011) 1:1–1:25, URL http://doi.acm.org/10.1145/2049662.2049663.